

ExBody2: Advanced Expressive Humanoid Whole-Body Control

Mazeyu Ji^{*1} Xuanbin Peng^{*1} Fangchen Liu² Jialong Li¹
Ge Yang³ Xuxin Cheng^{1,†} Xiaolong Wang^{1,4,†}
¹UC San Diego ²UC Berkeley ³MIT ⁴NVIDIA

^{*}Equal contribution [†]Equal advising

<https://exbody2.github.io>




Figure 1. Humanoid robot executing various expressive whole-body motions in the real world. The robot can (a) walk with a large stride from static standing, (b) dance along a long horizon (43 seconds) choreography, (c) crouch and stand up, (d) punch with different height configurations, (e,f) express various upper-body movements while maintaining balance.

Abstract

This paper enables real-world humanoid robots to maintain stability while performing expressive motions like humans do. We propose Advanced Expressive Whole-Body Control (ExBody2), a generalized whole-body tracking framework that can take any reference motion inputs and

control the humanoid to mimic the motion. The model is trained in simulation with Reinforcement Learning and then transferred to the real world. It decouples keypoint tracking with velocity control, and effectively leverages a privileged teacher policy to distill precise mimic skills into the

target student policy, which enables high-fidelity replication of dynamic movements such as running, crouching, dancing, and other challenging motions. We present a comprehensive qualitative and quantitative analysis of crucial design factors in the paper. We conduct our experiments on two humanoid platforms and demonstrate the superiority of our approach against state-of-the-arts, providing practical guidelines to pursue the extreme of whole-body control for humanoid robots. More results can be found at <https://exbody2.github.io>.

1. Introduction

Humanoid robots, with their similar morphology to humans, possess the potential to perform various tasks and motions that humans can accomplish in daily life. However, due to the high-dimensional state space and the complexity of control, it remains challenging to develop emergent human-like behaviors, limiting their real-world applications. With the growing availability of large-scale human motion datasets [4, 45], a practical approach to address this challenge is learning to replicate versatile movements by tracking and mimicking human motions [8, 20, 23, 24]. However, when considering the hardware, humanoid robots and humans are still quite different, which hinders the robot’s ability to replicate human motion completely. This poses a compelling research question: how can we pursue the expressive, human-like capabilities of humanoid robots while maintaining their stability and robustness, given their physical limitations?

In this paper, we introduce Advanced Expressive Whole-Body Control (ExBody2), an effective framework that maximizes the expressiveness of feasible whole-body motions on humanoid robots. The framework falls into a Sim2Real pipeline, where the policy will take a reference kinematic motion as input, and outputs the action that controls a real humanoid to conduct the motion in the real world. We train a single policy that generalizes across different input motions. We identify four technical designs to achieve this:

(i) *Constructing feasible and diverse training datasets.* Human motion datasets like AMASS [45] often contain complex movements beyond a robot’s physical capabilities, making tracking overly challenging and reducing performance. Some works address this by refining datasets. ExBody[8], for instance, filters out infeasible motions using language labels, though ambiguous descriptions (e.g., “dance”) can still include unsuitable movements. Other approaches, such as H2O[24] and OmniH2O [24], employ the SMPL model to simulate virtual humanoids and filter out complex motions. However, SMPL avatars can perform extreme actions that real robots cannot, creating a gap between simulation and real-world feasibility that can still impact training effectiveness. We systematically analyze dataset

difficulty and diversity, focusing on upper-body movement diversity for stability and tracking accuracy, and lower-body feasibility for robust training. Experiments demonstrate that diversity and feasibility in datasets are crucial for expressive and stable behaviors in both training and testing phases.

(ii) *Two-stage training.* ExBody2 adopts a two-stage teacher-student training framework as illustrated in Figure 2. We first train a teacher policy using a standard RL algorithm, PPO [60] with privileged information, which includes real root velocity, accurate positions for each body link, and other physical properties (e.g. frictions) in simulation. This results in an effective policy to accurately mimic versatile human motion. We then learn a deployable student policy by performing DAgger[59]-style distillation on the teacher policy. For the student policy, the privileged information is replaced by a series of historical observations.

(iii) *Local keypoint tracking strategy.* Previous whole-body tracking approaches, such as H2O [24] and OmniH2O [23], rely on *global tracking* of keypoint positions. This approach often leads to tracking failures in immediate next steps when robots struggle to align with current global keypoints, limiting their applications to highly stationary scenarios. In contrast, ExBody2 converts keypoints in the local frame, and decouples keypoint tracking from velocity control. To further improve the robustness of tracking, we periodically reset the global tracking keypoints to the local frame. This approach ensures sustained and robust whole-body tracking.

(iv) *Long-range motion with CVAE.* With enhanced whole-body tracking capabilities, our robot can now perform long-duration motion imitation. However, the motions in existing datasets are relatively short, presenting a challenge when aiming to enable continuous motion execution without manual resets. To overcome this, we train a Conditional Variational Autoencoder (CVAE) [64, 79] model to synthesize future motion. As shown in Figure 2, the CVAE takes on past motion information and generates future motion sequences, allowing the robot to execute complex and expressive movements seamlessly during deployment.

In the experiments, we evaluate four baseline methods on two robot platforms (Unitree G1 and H1) to assess the effects of different tracking methods, motion control strategies, and training techniques. Our approach demonstrates superior performance across full-body tracking accuracy, upper and lower body tracking, and velocity tracking, particularly excelling in velocity tracking and overall precision. Ablation studies further highlight that a curated dataset with moderate diversity and feasibility achieves optimal tracking performance, avoiding the drawbacks of overly simplistic or noisy datasets. By incorporating a teacher-student structure and an optimal history length, our approach achieves stable and efficient tracking in dynamic environments.

2. Related Work

Humanoid Whole-Body Control Whole-body control for humanoid robots remains a challenging problem due to the system’s high non-linearity. Traditional approaches predominantly rely on dynamics modeling and control [10–13, 26, 28, 33, 35, 49–51, 58, 73, 76]. Recent advances in reinforcement learning and sim-to-real transfer have demonstrated promising results in enabling complex whole-body skills for both quadrupedal and humanoid robots [1, 6, 7, 14, 16–19, 21, 29–32, 37, 39–42, 47, 48, 52, 56, 57, 61–63, 66, 72, 74, 75, 80]. Notably, works such as [8, 20, 24] have focused on humanoid whole-body control for expressive motions, with applications in manipulation and imitation learning. However, these approaches still exhibit limitations in expressiveness and maneuverability, highlighting the untapped potential of humanoid robots.

Robot Motion Imitation Robot motion imitation can be categorized into two primary areas: manipulation and expressiveness. For manipulation tasks, robots—often wheeled or tabletop—prioritize precise control over balancing and ground contact, making humanoid morphology unnecessary. Such robots typically utilize data from teleoperation [2, 3, 79] or human demonstrations [5, 34, 65, 70]. In contrast, expressive motion imitation focuses on learning life-like behaviors from human or animal motion capture data. This task is more challenging due to the need for fine-grained control of contacts and balance. While reinforcement learning has enabled physics-based character motion imitation in simulation [22, 43, 44, 54, 55, 67, 68, 71, 77], transferring such methods to real robots remains a significant challenge [8, 15, 16, 20, 23, 25, 53].

Human Motion Data Human motion capture datasets [4, 36, 45] offer a rich source of reference motions for training low-level policies. Additionally, recent advances in generative modeling enable the creation of multi-modal motion data using diffusion models conditioned on text inputs [69, 78] and motion variational autoencoders (VAEs) [43]. Our work leverages these kinematically feasible motion datasets to train physically plausible controllers, while the acquisition of motion data remains outside our scope.

3. ExBody2: Learning Expressive Humanoid Whole-Body Control

We propose Advanced Expressive Whole-Body Control (ExBody2), a simple and effective sim-to-real framework for expressive and robust whole-body control. As illustrated in Figure 2, ExBody2 consists of four main components: dataset curation, policy learning, motion synthesis, and real-world deployment. In this section, we will introduce these components of ExBody2 in detail.

3.1. Motion Dataset Curation

In the curation of our motion datasets, significant emphasis is placed on the detailed analysis and selection of actions, particularly distinguishing between the capabilities required for upper versus lower body movements. This strategic approach aims to optimize the diversity of movements within feasible limits for robotic implementations.

Upper Body Actions Our dataset includes a wide array of upper body movements. This variety is crucial to challenge and enhance the robot’s adaptability and readiness for real-world applications, where flexibility in response to unexpected scenarios is vital.

Lower Body Actions The selection of lower body movements is approached with greater conservatism due to the mechanical limitations and stability requirements of robots. Basic locomotive actions such as walking and subtle stance adjustments are included. However, highly dynamic movements such as jumping or complex rotational movements are carefully assessed. Training on such infeasible actions would introduce noise into the learning process, thus reducing its effectiveness.

Balancing Act and Dataset Integrity The process of curating the datasets involves a delicate balance between challenging the robotic capabilities and maintaining actions within a feasible spectrum. Overly simplistic tasks could limit the training policy’s ability to generalize to new situations, while overly complex tasks might exceed the robot’s operational capabilities, leading to ineffective learning outcomes. Part of our dataset preparation, therefore, includes the exclusion or modification of entries that featured complex lower body movements beyond the robot’s capabilities. This step is crucial to ensure the datasets are rich in actionable content without overwhelming the robot’s learning algorithms with impractical tasks.

Through meticulous planning, we curate distinct datasets that are optimally balanced for robotic training. Extensive experimental validation will demonstrate the superiority of our dataset selection approach in Section 4.

3.2. Policy Learning

ExBody2 aims at tracking a target motion more expressively in the whole body. To this end, ExBody2 adopts an effective two-stage teacher-student training procedure as in [37, 38]. Specifically, the oracle teacher policy is first trained with an off-the-shelf reinforcement learning (RL) algorithm, PPO [60], with privileged information that can be obtained only in simulators. For the second stage, we replace the privileged information with observations which are aligned with the real world, and distill the teacher policy into a deployable student policy. We train our policies using IsaacGym [46] with efficient parallel simulation.

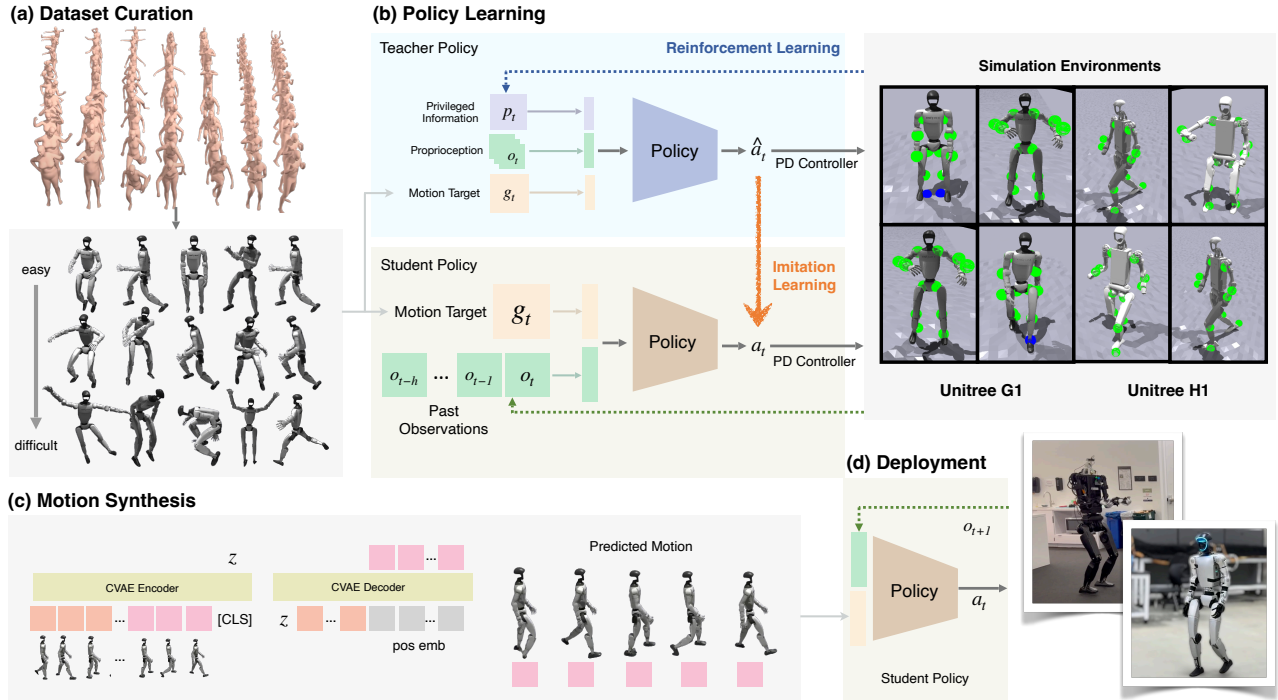


Figure 2. ExBody2’s framework. (a) ExBody2 presents a curated high-quality motion dataset, with highly diverse and feasible motions. (b) ExBody2 adopts a teacher-student learning framework on Unitree G1 and H1 whole-body control tasks. (c) ExBody2 employs a transformer-based CVAE for continuous motion synthesis. (d) By tracking the generated motion from CVAE, ExBody2 can be deployed seamlessly in the real world to mimic sustained motions expressively.

3.2.1. Teacher Policy Training

We can formulate the humanoid motion control problem as a *Markov Decision Process* (MDP). The state space \mathcal{S} contains privileged observation \mathcal{X} , proprioceptive states \mathcal{O} and motion tracking target \mathcal{G} . A policy $\hat{\pi}$ takes $\{p_t, o_t, g_t\}$ as input, and outputs action \hat{a}_t , as illustrated in Figure 2, the teacher policy. The predicted action $\hat{a}_t \in R^{23}$ is the target joint positions of joint proportional derivative (PD) controllers. We use off-the-shelf PPO [60] algorithm to maximize expectation of the accumulated future rewards $E_{\hat{\pi}}[\sum_{t=0}^T \gamma^t \mathcal{R}(s_t, \hat{a}_t)]$, which encourages tracking the demonstrations with robust behavior. The predicted $\hat{a}_t \in R^{23}$, which is the target position of joint proportional derivative (PD) controllers.

Privileged Information The privileged information p_t contains some ground-truth states of the humanoid robot and the environment, which can only be observed in simulators. It contains the ground-truth root velocity, real body links’ positions, and physical properties (e.g. friction coefficients, motor strength). Privileged information can significantly improve the sample efficiency of RL algorithms, which is often leveraged to obtain a high-performing teacher policy.

Motion Tracking Target Similar to Exbody [9], ExBody2 learns a policy that can be controlled by the joystick commands (e.g. the linear velocity and body pose) when

accurately tracking a whole-body motion. The motion tracking target consists of two components, which are (1) the desired joints and 3D keypoints in both the *upper and lower body* and (2) target root velocity and root pose. For the full information about the privileged information, motion tracking information, and proprioceptive observations for the teacher policy, please refer to the supplementary materials.

Reward Design Our reward function is meticulously constructed to enhance the performance and realism of the humanoid robot’s motion. The primary components of the reward include tracking the velocity, direction, and orientation of the root, alongside precise tracking of keypoints and joint positions. Additionally, we incorporate several regularization terms designed to boost the robot’s stability and enhance the transferability from simulation to real-world applications. The main elements of our tracking reward are detailed in Table 1, while supplementary rewards aimed at improving stability and sim2real capabilities would be discussed extensively in the supplementary materials.

3.2.2. Student Policy Training

In this stage, we remove the privileged information, and use longer history observation to train a student policy. As shown in Figure 2, the student policy encodes a series of

past observations $o_{t-H:t}$ together with the encoded g_t to get the predicted $a_t \sim \pi(\cdot | o_{t-H:t}, g_t)$. We supervise π using the teacher’s action $\hat{a}_t \sim \hat{\pi}(\cdot | o_t, g_t)$ with an MSE loss.

$$l = \|a_t - \hat{a}_t\|^2$$

To train the student, we adopt the strategy used in DAgger [59], we roll out the student policy π in the simulation environment to generate training data. For each visited state, the teacher policy $\hat{\pi}$ computes the oracle action as the supervision signal. We proceed to refine the policy π by iteratively minimizing the loss l on the accumulated data. The training of $\hat{\pi}$ continues through successive rollouts until the loss l reaches convergence. A critical aspect of training the student policy is preserving a sufficiently long sequence of historical observations. Detailed results and further analysis are elaborated in Section 4.

Term	Expression	Weight
Expression Goal G^e		
DoF Position	$\exp(-0.7 \mathbf{q}_{\text{ref}} - \mathbf{q})$	3.0
Keypoint Position	$\exp(- \mathbf{p}_{\text{ref}} - \mathbf{p})$	2.0
Root Movement Goal G^m		
Linear Velocity	$\exp(-4.0 \mathbf{v}_{\text{ref}} - \mathbf{v})$	6.0
Velocity Direction	$\exp(-4.0 \cos(\mathbf{v}_{\text{ref}}, \mathbf{v}))$	6.0
Roll & Pitch	$\exp(- \boldsymbol{\Omega}_{\text{ref}}^{\phi\theta} - \boldsymbol{\Omega}^{\phi\theta})$	1.0
Yaw	$\exp(- \Delta y)$	1.0

Table 1. Rewards Specification for ExBody2.

3.2.3. Decomposed Tracking Strategy

Motion tracking comprises two objectives: tracking DoF (joint) positions and keypoint (body keypoint) positions. Keypoint tracking usually plays a crucial role in tracking motions, as joint DoF errors can propagate to the whole body, while keypoint tracking is directly applied to the body. Existing work like H2O, OmniH2O [23, 24] learns to follow the trajectory of global keypoints. However, this global tracking strategy usually results in suboptimal or failed tracking behavior, as global keypoints may drift through time, resulting in cumulative errors that eventually hinder learning. To address this, we map global keypoints to the robot’s current coordinate frame, and instead utilize velocity-based global tracking. The coordination of velocity and motion allows tracking completion with maximal expressiveness, even if slight positional deviations arise. Moreover, to further enhance the robot’s capabilities in following challenging keypoint motions, we allow a small global drift of keypoints during training stage and periodically correct them to the robot’s current coordinate frame. During deployment, we strictly employ local keypoint tracking with velocity-decomposed control.

3.3. Continual Motion Synthesis

The existing motions are usually short in length, which limits humanoid robots from performing interesting and extended behaviors continuously. To overcome this, we train a Conditional Variational Autoencoder (CVAE) [64, 79] to synthesize future movements that can be directly employed by the whole-body tracking policy, allowing the robot to execute complex and expressive movements seamlessly during deployment.

As shown in Figure 2 (c), the CVAE takes on past M -step motions $m_{t-M:t}$, and synthesizes future H -step motions $m_{t+1:t+H+1}$ autoregressively. m_t includes the current DoF positions, root pose, velocity, and angular velocity. Specifically, we utilize a transformer-based CVAE architecture, similar to the architecture proposed in [79]. We first separately tokenize DoF positions and root information through two separate MLPs, and get tokenized motion information. We compute the posterior $q_\phi(m_{t-M:t}, m_{t+1:t+H+1})$ through a transformer encoder network, using the features from the $[CLS]$ token. For notational simplicity and consistency with standard CVAE, we denote the condition $c = m_{t-M:t}$ to be the history context, and $x = m_{t+1:t+H+1}$ to be the prediction target. During training time, the latent variable z is sampled from $\mathcal{N}(\mu(c, x), \sigma(c, x)^2)$. A transformer decoder will take z with $c = m_{t-M:t}$, and positional embeddings to predict H -step future motions \hat{x} . The CVAE loss consists of the reconstruction loss $l_{recon} = \|x - \hat{x}\|^2$ and the KL divergence loss:

$$l_{KL} = \frac{1}{2} \sum_{i=1}^d (\sigma_i^2(x, c) + \mu_i^2(x, c) - 1 - \log(\sigma_i^2(x, c)))$$

where d is the dimensionality of the latent space. To improve the smoothness of the predicted motion, we also apply a smoothness loss:

$$l_{smooth} = \|\hat{m}_{t+1} - m_t\|^2 + \sum_{i=t+1}^{t+H} \|\hat{m}_{i+1} - \hat{m}_i\|^2$$

The total loss for cVAE training is: $l_{motion} = l_{recon} + \alpha l_{KL} + \beta l_{smooth}$ where we set $\alpha = \beta = 0.5$ in our experiments. During inference, z is set to the mean of the prior (i.e. zero). We adopt a similar temporal ensemble strategy as in Zhao et al. [79] to get predicted future motions. Further details regarding the encoder and decoder architecture of the CVAE are in the supplementary materials.

4. Experiments

4.1. Experimental Setup

We conduct our experiments in IsaacGym [46] simulator across two robot platforms (Unitree G1 and H1).

Method	Sim2Real	$E_{vel} \downarrow$	$E_{mpkpe} \downarrow$	$E_{mpkpe}^{upper} \downarrow$	$E_{mpkpe}^{lower} \downarrow$	$E_{mpjpe} \downarrow$	$E_{mpjpe}^{upper} \downarrow$	$E_{mpjpe}^{lower} \downarrow$
Robot G1								
Exbody [†]	✓	0.2950	0.0914	0.0830	0.0997	0.1145	0.0989	0.1359
Exbody	✓	0.3083	0.1134	0.0924	0.1343	0.1660	0.0766	0.2891
OmniH2O*	✓	0.2429	0.0959	0.0878	0.1041	0.1103	0.0913	0.1364
Exbody2	✓	0.1891	0.0854	0.0781	0.0928	0.0938	0.0784	0.1149
Robot H1								
Exbody [†]	✗	0.6377	0.1950	0.2055	0.1827	0.3029	0.3028	0.3029
Exbody	✓	0.5964	0.1941	0.1948	0.1916	0.3251	0.2784	0.3775
OmniH2O*	✓	0.4254	0.1540	0.1563	0.1509	0.2184	0.2135	0.2240
Exbody2	✓	0.3016	0.1290	0.1317	0.1247	0.1763	0.1722	0.1809

Table 2. Simulation motion imitation evaluation of **Exbody2** and baselines on dataset \mathcal{D}_{250} for robots G1 and H1. Statistically significant results are highlighted in bold across 5 random seeds.

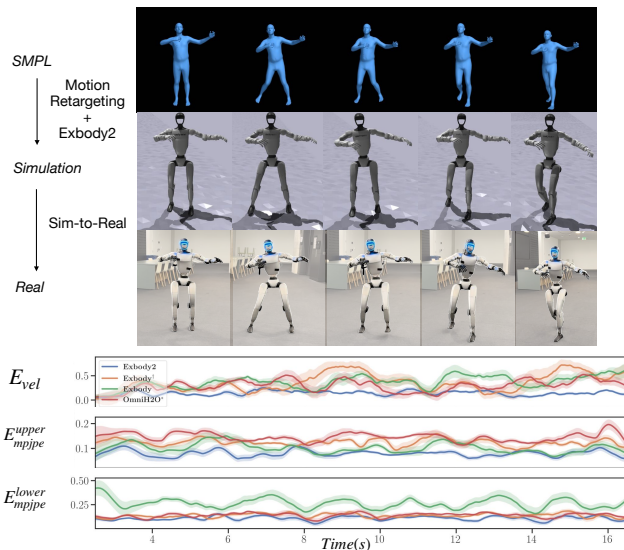


Figure 3. A sequence of a robot performing the Cha-Cha dance. From top to bottom: the reference motion represented by an avatar, our algorithm’s performance in the simulation, and its performance on a real robot. The bottom three rows show the per-frame errors: velocity tracking error, upper-body joint DOF error, and lower-body DOF error, with the blue curve representing **ExBody2**, orange for **Exbody[†]**, green for **Exbody**, and red for **OmniH2O***.

Baselines. To examine the effectiveness of varying tracking methods, motion control strategies, and training techniques, we evaluate four baselines using a high-quality, curated dataset of 250 motion samples. This dataset is carefully selected for optimal training suitability, as will be shown in subsequent experiments.

- **Exbody** [9]: This method only tracks the upper body movements from the human data, while tracking the root motion of the lower body without explicitly following step patterns and focusing on partial body tracking.
- **Exbody[†]**: The whole-body control version of **Exbody**, where the full-body movements are tracked based on human data. This setup enables comprehensive human motion imitation, attempting to match the entire body pos-

ture to the reference data.

- **OmniH2O***: Our reimplementation of the OmniH2O [23], using global keypoints tracking and the same observation space as described in the original paper. During deployment, we adapt OmniH2O to our local tracking evaluation for fair comparison.
- **ExBody2**: Our method utilizes local keypoint tracking. It incorporates various training techniques to enhance the overall motion fidelity and sim-to-real transfer.

Metrics. We evaluate the policy’s performance using several metrics calculated across all motion sequences from the dataset. The *mean linear velocity error* E_{vel} (m/s) measures the error between the robot’s root linear velocity and that of the demonstration, indicating the policy’s ability to track velocity. We calculate tracking error in terms of keybody positions and joint angles. The *Mean Per Keybody Position Error* (MPKPE) E_{mpkpe} (m) evaluates the overall keypoint position tracking ability. For more detailed analysis, we report the MPKPE for the upper body E_{mpkpe}^{upper} (m) and the lower body E_{mpkpe}^{lower} (m), assessing the keypoint position tracking of the upper and lower body, respectively. Similarly, the *Mean Per Joint Position Error* (MPJPE) E_{mpjpe} (rad) measures the joint tracking ability. We also report the MPJPE for the upper body E_{mpjpe}^{upper} (rad) and the lower body E_{mpjpe}^{lower} (rad) to evaluate tracking performance in different body regions.

More details about environment designs and baseline implementations can be found in the supplementary material.

4.2. Main Results

Table 2 and Figure 3 present the results of different algorithms on two robot platforms, Unitree G1 and H1, where **ExBody2** outperforms all other methods in all the tracking metrics. Our approach maintains consistent lower tracking error across the entire time frame with minimal fluctuations, indicating a stable and precise tracking performance. To be more precise, we evaluate the following metrics:

1. **Upper Body Tracking Accuracy:** Our algorithm per-

Training Dataset	In dist.	Metrics						
		$E_{vel} \downarrow$	$E_{mpkpe} \downarrow$	$E_{mpkpe}^{upper} \downarrow$	$E_{mpkpe}^{lower} \downarrow$	$E_{mpjpe} \downarrow$	$E_{mpjpe}^{upper} \downarrow$	$E_{mpjpe}^{lower} \downarrow$
(a) Eval. on \mathcal{D}_{50}								
\mathcal{D}_{50}	✓	0.1375	0.0627	0.0571	0.0682	0.0753	0.0626	0.0928
\mathcal{D}_{250}	✓	0.1454	0.0669	0.0600	0.0738	0.0870	0.0689	0.1119
\mathcal{D}_{500}	✓	0.1593	0.0740	0.0660	0.0820	0.1051	0.0865	0.1307
\mathcal{D}_{CMU}	✓	0.1543	0.0767	0.0649	0.0885	0.1099	0.0854	0.1437
(b) Eval. on \mathcal{D}_{CMU}								
\mathcal{D}_{50}	✗	0.2621	0.1076	0.1071	0.1078	0.1343	0.1321	0.1373
\mathcal{D}_{250}	✗	0.1901	0.1048	0.1009	0.1084	0.1119	0.0998	0.1286
\mathcal{D}_{500}	✗	0.1966	0.1083	0.1044	0.1124	0.1252	0.1106	0.1452
\mathcal{D}_{CMU}	✓	0.1788	0.1106	0.1047	0.1171	0.1308	0.1130	0.1552
(c) Eval. on \mathcal{D}_{ACCAD}								
\mathcal{D}_{50}	✗	0.4226	0.1277	0.1210	0.1330	0.1720	0.1618	0.1861
\mathcal{D}_{250}	✗	0.3533	0.1234	0.1141	0.1315	0.1421	0.1223	0.1692
\mathcal{D}_{500}	✗	0.3581	0.1280	0.1198	0.1355	0.1586	0.1400	0.1842
\mathcal{D}_{CMU}	✗	0.3452	0.1267	0.1146	0.1381	0.1780	0.1635	0.1979

Table 3. Dataset Ablation Study: Evaluation on \mathcal{D}_{50} , \mathcal{D}_{CMU} , and \mathcal{D}_{ACCAD} datasets with models trained on various datasets. Statistically significant results are highlighted in bold across 5 random seeds.

forms slightly lower than *Exbody* in upper body tracking accuracy. However, *Exbody* sacrifices lower body tracking to maintain balance, whereas our algorithm can perform well with balanced tracking accuracy for both upper body and lower body.

- Lower Body Tracking Accuracy:** Our algorithm demonstrates higher accuracy in lower body tracking compared to other algorithms for both G1 and H1.
- Velocity Tracking Accuracy:** Our algorithm also excels in velocity tracking accuracy, outperforming other methods. This improvement is attributed to the teacher-student training paradigm, where the distillation from privileged information to historical observation results in a student policy with better velocity-tracking.

Overall, our algorithm achieves significant improvements in full-body tracking accuracy for both upper and lower body, and velocity tracking accuracy compared to the baseline algorithms, demonstrating stable and effective tracking performance in dynamic environments.

4.3. Ablation on Datasets

We manually design four datasets of varying sizes, where the largest being the complete CMU AMASS dataset, containing in total 1,919 motions. The remaining datasets, with sizes 50, 250, and 500, are subsets of the CMU dataset, each constructed with different levels of action diversity:

- **50-action dataset (\mathcal{D}_{50}):** This subset includes only fun-

damental static actions, such as standing and walking. Both upper and lower limb actions lack diversity, primarily covering basic and easily feasible motions.

- **250-action dataset (\mathcal{D}_{250}):** In addition to basic actions, this subset introduces more diverse upper limb actions and a variety of lower limb actions, such as running and simple dance movements. These actions are feasible for robots, albeit moderately challenging.
- **500-action dataset (\mathcal{D}_{500}):** This subset includes increasingly complex lower limb actions, such as single-leg jumps and intricate dance moves, rendering the dataset somewhat noisier with less feasible motions.
- **Full CMU dataset (\mathcal{D}_{CMU}):** The complete CMU dataset comprises 1,919 actions, including extreme movements like push-ups, rolling on the ground, and somersaults. This dataset is unfiltered and introduces significant noise.

We train ExBody2 on the above four datasets, and conduct three sets of evaluations on the G1 robot, shown in Table 3:

- Evaluation on \mathcal{D}_{50} :** Policies trained on \mathcal{D}_{50} achieve the highest tracking accuracy for in-distribution actions, as reflected in metrics across all categories. This suggests that additional data does not necessarily benefit in-distribution tasks. While the policy trained on \mathcal{D}_{250} performs similarly to \mathcal{D}_{50} , policies trained on \mathcal{D}_{500} and \mathcal{D}_{CMU} exhibit substantial drops in tracking accuracy.
- Evaluation on \mathcal{D}_{CMU} :** Policies trained on \mathcal{D}_{250} achieve the best performance on \mathcal{D}_{CMU} , surpassing those

Method	$E_{vel} \downarrow$	$E_{mpkpe} \downarrow$	$E_{mpkpe}^{upper} \downarrow$	$E_{mpkpe}^{lower} \downarrow$	$E_{mpipe} \downarrow$	$E_{mpipe}^{upper} \downarrow$	$E_{mpipe}^{lower} \downarrow$
(a) History Length Ablation							
Exbody2-History25 (Ours)	0.1891	0.0854	0.0781	0.0928	0.0938	0.0784	0.1149
Exbody2-History0	0.3101	0.0896	0.0812	0.0980	0.0972	0.0805	0.1201
Exbody2-History10	0.1934	0.0860	0.0792	0.0927	0.0935	0.0777	0.1152
Exbody2-History50	0.1944	0.0871	0.0797	0.0945	0.0962	0.0791	0.1198
Exbody2-History100	0.2119	0.0854	0.0778	0.0929	0.0960	0.0787	0.1198
(b) Reset Global Tracking Ablation							
Exbody2(Ours)	0.1891	0.0854	0.0781	0.0928	0.0938	0.0784	0.1149
Exbody2-w/o-Reset	0.2545	0.0871	0.0807	0.0935	0.0919	0.0752	0.1149
(c) DAgger Ablation							
Exbody2(Ours)	0.1891	0.0854	0.0781	0.0928	0.0938	0.0784	0.1149
Exbody2-w/o-DAgger	0.3115	0.0925	0.0843	0.1006	0.1152	0.0996	0.1366

Table 4. Self Ablation Study: Evaluation of different configurations of our method on dataset D_{250} . The table is divided into three parts: (a) History Length Ablation, (b) Reset Global Tracking Ablation, and (c) DAgger Ablation. The best results in each column are highlighted.

trained on both D_{500} and the full D_{CMU} dataset. Due to the limited diversity of the D_{50} dataset, especially in upper limb movements, the D_{50} -trained policy struggles to maintain high accuracy for out-of-distribution actions. Unexpectedly, the D_{250} -trained policy generalizes better than those trained on D_{500} and D_{CMU} . This result underscores that noisy datasets degrade policy performance, as the policy may expend unnecessary effort on tracking infeasible actions, lowering the accuracy of feasible actions.

- 3. Evaluation on D_{ACCAD} :** This experiment further emphasizes the importance of clean datasets. Here, the ACCAD dataset (D_{ACCAD}) consists of actions that are entirely not in the training data. Policy trained on D_{250} outperforms the others, achieving the best tracking accuracy. Additionally, the D_{250} , D_{500} , and D_{CMU} -trained policies perform relatively well in velocity tracking. However, the D_{50} -trained policy suffers from substantial tracking errors, suggesting the limitations of a small, simple dataset in handling unseen data.

In conclusion, the D_{250} dataset yields the best performance, providing a balance between diversity and feasibility. It includes a rich variety of upper limb actions and moderately challenging lower limb actions that remain feasible for robots. This dataset avoids the drawbacks of overly simplistic datasets with insufficient diversity and noisy datasets that degrade tracking performance.

4.4. Ablation on Policy Training

Finally, we ablate the design choices in ExBody2, including the effectiveness of the teacher-student setup, the optimal history length to train the policy, and the mechanism of using delayed keypoints reset, as shown in Table 4. The experiments are conducted on the Unitree G1 robot.

We first examine the impact of teacher-student training.

As shown in Table 4 (c), the tracking accuracy drops significantly without teacher-student training. This is largely due to the absence of privileged velocity guidance, making it difficult for the single-stage RL policy to learn velocity directly from history. Consequently, it could not keep up with the desired movements, leading to worse tracking precision.

We then test student policies trained with different history lengths in Table 4 (a). When no extra history is used, the policy struggles to learn effectively. Among the non-zero history lengths, most policies perform similarly while the history length of 25 yields the best results, which is used by us in the main experiments. Longer history lengths increase the difficulty of fitting the privileged information, ultimately reducing tracking performance.

Lastly, we evaluate the impact of adding small drifting by using a delayed global keypoint reset strategy in Table 4 (b). We allow the keypoints to move in the global frame for short periods before resetting to the robot’s local frame. This effectively serves as a data augmentation strategy, encouraging the robot to track keypoints with more significant drift. This design notably improves velocity tracking, enabling the robot to use absolute position adjustments to compensate for discrepancies when tracking velocities.

5. Conclusion

In this work, we develop ExBody2, a novel algorithm that achieves stable and expressive whole-body motion for humanoid robots by utilizing a curated dataset and a decoupled whole-body tracking framework. Our approach separates key body tracking from velocity control, enabling precise and stable motion replication. While we have identified the characteristics of datasets best suited for humanoid robot training, an automated method to select high-quality datasets remains an open challenge. Addressing this limitation will be a focus of our future work.

References

- [1] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning*, pages 403–415. PMLR, 2023. 3
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 3
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. 3
- [4] Carnegie Mellon University. Carnegie-Mellon mocap database. <http://mocap.cs.cmu.edu/>, 2007. [Online]. 2, 3
- [5] Sirui Chen, Chen Wang, Kaden Nguyen, Li Fei-Fei, and C Karen Liu. Arcap: Collecting high-quality human demonstrations for robot learning with augmented reality feedback. *arXiv preprint arXiv:2410.08464*, 2024. 3
- [6] Xuxin Cheng, Ashish Kumar, and Deepak Pathak. Legs as manipulator: Pushing quadrupedal agility beyond locomotion. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023. 3
- [7] Xuxin Cheng, Kexin Shi, Ananye Agarwal, and Deepak Pathak. Extreme parkour with legged robots. *arXiv preprint arXiv:2309.14341*, 2023. 3
- [8] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots, 2024. 2, 3
- [9] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024. 4, 6, 12
- [10] Matthew Chignoli, Donghyun Kim, Elijah Stanger-Jones, and Sangbae Kim. The mit humanoid robot: Design, motion planning, and control for acrobatic behaviors. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 1–8. IEEE, 2021. 3
- [11] Antonin Dallard, Mehdi Benallegue, Fumio Kanehiro, and Abderrahmane Kheddar. Synchronized human-humanoid motion imitation. *IEEE Robotics and Automation Letters*, 8(7):4155–4162, 2023.
- [12] Behzad Dariush, Michael Gienger, Bing Jian, Christian Goller, and Kikuo Fujimura. Whole body humanoid control from human motion descriptors. In *2008 IEEE International Conference on Robotics and Automation*, pages 2677–2684. IEEE, 2008.
- [13] Kouros Darvish, Yeshasvi Tirupachuri, Giulio Romualdi, Lorenzo Rapetti, Diego Ferigo, Francisco Javier Andrade Chavez, and Daniele Pucci. Whole-body geometric retargeting for humanoid robots. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, pages 679–686, 2019. 3
- [14] Helei Duan, Bikram Pandit, Mohitvishnu S Gadde, Bart Jaap van Marum, Jeremy Dao, Chanho Kim, and Alan Fern. Learning vision-based bipedal locomotion for challenging terrain. *arXiv preprint arXiv:2309.14594*, 2023. 3
- [15] Pranay Dugar, Aayam Shrestha, Fangzhou Yu, Bart van Marum, and Alan Fern. Learning multi-modal whole-body control for real-world humanoid robots, 2024. 3
- [16] Alejandro Escontrela, Xue Bin Peng, Wenhao Yu, Tingnan Zhang, Atil Iscen, Ken Goldberg, and Pieter Abbeel. Adversarial motion priors make good substitutes for complex reward functions. 2022 iee. In *International Conference on Intelligent Robots and Systems (IROS)*, 2022. 3
- [17] Jiawei Fu, Yunlong Song, Yan Wu, Fisher Yu, and Davide Scaramuzza. Learning deep sensorimotor policies for vision-based autonomous drone racing, 2022.
- [18] Zipeng Fu, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Minimizing energy consumption leads to the emergence of gaits in legged robots. *Conference on Robot Learning (CoRL)*, 2021.
- [19] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: learning a unified policy for manipulation and locomotion. In *Conference on Robot Learning*, pages 138–149. PMLR, 2023. 3
- [20] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans, 2024. 2, 3
- [21] Yuni Fuchioka, Zhaoming Xie, and Michiel Van de Panne. Opt-mimic: Imitation of optimized trajectories for dynamic quadruped behaviors. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5092–5098. IEEE, 2023. 3
- [22] Mohamed Hassan, Yunrong Guo, Tingwu Wang, Michael Black, Sanja Fidler, and Xue Bin Peng. Synthesizing physical character-scene interactions. 2023. 3
- [23] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning, 2024. 2, 3, 5, 6, 12
- [24] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. *arXiv preprint arXiv:2403.04436*, 2024. 2, 3, 5
- [25] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Changliu Liu, Guanya Shi, Xiaolong Wang, Linxi Fan, and Yuke Zhu. Hover: Versatile neural whole-body controller for humanoid robots. *arXiv preprint arXiv:2410.21229*, 2024. 3
- [26] Kazuo Hirai, Masato Hirose, Yuji Haikawa, and Toru Takenaka. The development of honda humanoid robot. In *Proceedings. 1998 IEEE international conference on robotics and automation (Cat. No. 98CH36146)*, pages 1321–1326. IEEE, 1998. 3
- [27] Albert S. Huang, Edwin Olson, and David C. Moore. Lcm: Lightweight communications and marshalling. *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4057–4062, 2010. 12

- [28] Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, et al. Anymal-a highly mobile and dynamic quadrupedal robot. In *IROS*, 2016. 3
- [29] Hiroshi Ito, Kenjiro Yamamoto, Hiroki Mori, and Tetsuya Ogata. Efficient multitask learning with an embodied predictive model for door opening and entry with whole-body control. *Science Robotics*, 7(65):eaax8177, 2022. 3
- [30] Seunghun Jeon, Moonkyu Jung, Suyoung Choi, Beomjoon Kim, and Jemin Hwangbo. Learning whole-body manipulation for quadrupedal robot. *arXiv preprint arXiv:2308.16820*, 2023.
- [31] Yandong Ji, Zhongyu Li, Yanan Sun, Xue Bin Peng, Sergey Levine, Glen Berseth, and Koushil Sreenath. Hierarchical reinforcement learning for precise soccer shooting skills using a quadrupedal robot. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1479–1486. IEEE, 2022.
- [32] Yandong Ji, Gabriel B Margolis, and Pulkit Agrawal. Dribblebot: Dynamic legged manipulation in the wild. *arXiv preprint arXiv:2304.01159*, 2023. 3
- [33] Shuuji Kajita, Fumio Kanehiro, Kenji Kaneko, Kazuhito Yokoi, and Hirohisa Hirukawa. The 3d linear inverted pendulum mode: A simple modeling for a biped walking pattern generation. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, pages 239–246. IEEE, 2001. 3
- [34] Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranay Mathur, Shuo Cheng, Chen Wang, Judy Hoffman, and Danfei Xu. Egomimic: Scaling imitation learning via egocentric video, 2024. 3
- [35] Ichiro Kato. Development of wabot 1. *Biomechanism*, 2: 173–214, 1973. 3
- [36] Franziska Krebs, Andre Meixner, Isabel Patzer, and Tamim Asfour. The kit bimanual manipulation dataset. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 499–506, 2021. 3
- [37] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021. 3
- [38] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020. 3
- [39] Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2811–2817. IEEE, 2021. 3
- [40] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Robust and versatile bipedal jumping control through multi-task reinforcement learning. *arXiv preprint arXiv:2302.09450*, 2023.
- [41] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. *arXiv preprint arXiv:2401.16889*, 2024.
- [42] Qiayuan Liao, Bike Zhang, Xuanyu Huang, Xiaoyu Huang, Zhongyu Li, and Koushil Sreenath. Berkeley humanoid: A research platform for learning-based control. *arXiv preprint arXiv:2407.21781*, 2024. 3
- [43] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)*, 39(4):40–1, 2020. 3
- [44] Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris M. Kitani, and Weipeng Xu. Universal humanoid motion representations for physics-based control. In *The Twelfth International Conference on Learning Representations*, 2024. 3
- [45] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. Amass: Archive of motion capture as surface shapes. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3
- [46] Viktor Makovychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021. 3, 5
- [47] Gabriel B Margolis, Tao Chen, Kartik Paigwar, Xiang Fu, Donghyun Kim, Sangbae Kim, and Pulkit Agrawal. Learning to jump from pixels. *arXiv preprint arXiv:2110.15344*, 2021. 3
- [48] Gabriel B Margolis, Ge Yang, Kartik Paigwar, Tao Chen, and Pulkit Agrawal. Rapid locomotion via reinforcement learning. *arXiv preprint arXiv:2205.02824*, 2022. 3
- [49] Hirofumi Miura and Isao Shimoyama. Dynamic walk of a biped. *IJRR*, 1984. 3
- [50] Federico L Moro and Luis Sentis. Whole-body control of humanoid robots. *Humanoid Robotics: A reference*, Springer, Dordrecht, 2019.
- [51] Luigi Penco, Nicola Scianca, Valerio Modugno, Leonardo Lanari, Giuseppe Oriolo, and Serena Ivaldi. A multimode teleoperation framework for humanoid loco-manipulation: An application for the icub robot. *IEEE Robotics and Automation Magazine*, 26(4):73–82, 2019. 3
- [52] Xue Bin Peng, Erwin Coumans, Tingnan Zhang, Tsang-Wei Lee, Jie Tan, and Sergey Levine. Learning agile robotic locomotion skills by imitating animals. 2020. 3
- [53] Xue Bin Peng, Erwin Coumans, Tingnan Zhang, Tsang-Wei Edward Lee, Jie Tan, and Sergey Levine. Learning agile robotic locomotion skills by imitating animals. In *Robotics: Science and Systems*, 2020. 3
- [54] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (ToG)*, 40(4):1–20, 2021. 3
- [55] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Trans. Graph.*, 41(4), 2022. 3

- [56] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *arXiv:2303.03381*, 2023. 3
- [57] Ilija Radosavovic, Bike Zhang, Baifeng Shi, Jathushan Rajasegaran, Sarthak Kamat, Trevor Darrell, Koushil Sreenath, and Jitendra Malik. Humanoid locomotion as next token prediction. *arXiv:2402.19469*, 2024. 3
- [58] Joao Ramos and Sangbae Kim. Dynamic locomotion synchronization of bipedal robot and human operator via bilateral feedback teleoperation. *Science Robotics*, 4(35): eaav4282, 2019. 3
- [59] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011. 2, 5, 13
- [60] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 2, 3, 4, 13
- [61] Clemens Schwarke, Victor Klemm, Matthijs Van der Boon, Marko Bjelonic, and Marco Hutter. Curiosity-driven learning of joint locomotion and manipulation tasks. In *Proceedings of The 7th Conference on Robot Learning*, pages 2594–2610. PMLR, 2023. 3
- [62] Mingyo Seo, Steve Han, Kyutae Sim, Seung Hyeon Bang, Carlos Gonzalez, Luis Sentis, and Yuke Zhu. Deep imitation learning for humanoid loco-manipulation through human teleoperation. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, pages 1–8. IEEE, 2023.
- [63] Jonah Siekmann, Kevin Green, John Warila, Alan Fern, and Jonathan Hurst. Blind bipedal stair traversal via sim-to-real reinforcement learning. *arXiv preprint arXiv:2105.08328*, 2021. 3
- [64] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015. 2, 5
- [65] Mohan Kumar Srirama, Sudeep Dasari, Shikhar Bahl, and Abhinav Gupta. Hrp: Human affordances for robotic pre-training. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024. 3
- [66] Annan Tang, Takuma Hiraoka, Naoki Hiraoka, Fan Shi, Kento Kawaharazuka, Kunio Kojima, Kei Okada, and Masayuki Inaba. Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation. *arXiv preprint arXiv:2309.14225*, 2023. 3
- [67] Chen Tessler, Yoni Kasten, Yunrong Guo, Shie Mannor, Gal Chechik, and Xue Bin Peng. Calm: Conditional adversarial latent models for directable virtual characters. In *ACM SIGGRAPH 2023 Conference Proceedings*, New York, NY, USA, 2023. Association for Computing Machinery. 3
- [68] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 2024. 3
- [69] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [70] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C. Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788*, 2024. 3
- [71] Jiashun Wang, Jessica Hodgins, and Jungdam Won. Strategy and skill learning for physics-based table tennis animation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 3
- [72] Yikai Wang, Zheyuan Jiang, and Jianyu Chen. Amp in the wild: Learning robust, agile, natural legged locomotion skills. *arXiv preprint arXiv:2304.10888*, 2023. 3
- [73] Eric R Westervelt, Jessy W Grizzle, and Daniel E Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE transactions on automatic control*, 48(1):42–56, 2003. 3
- [74] Ruihan Yang, Zhuoqun Chen, Jianhan Ma, Chongyi Zheng, Yiyu Chen, Quan Nguyen, and Xiaolong Wang. Generalized animal imitator: Agile locomotion with versatile motion prior. *arXiv preprint arXiv:2310.01408*, 2023. 3
- [75] Ruihan Yang, Ge Yang, and Xiaolong Wang. Neural volumetric memory for visual locomotion control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1430–1440, 2023. 3
- [76] KangKang Yin, Kevin Loken, and Michiel Van de Panne. Simbicon: Simple biped locomotion control. *ACM Transactions on Graphics*, 2007. 3
- [77] Haotian Zhang, Ye Yuan, Viktor Makoviychuk, Yunrong Guo, Sanja Fidler, Xue Bin Peng, and Kayvon Fatahalian. Learning physically simulated tennis skills from broadcast videos. *ACM Trans. Graph.*, 42(4), 2023. 3
- [78] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 3
- [79] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023. 2, 3, 5, 12
- [80] Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher Atkeson, Sören Schwertfeger, Chelsea Finn, and Hang Zhao. Robot parkour learning. In *Conference on Robot Learning (CoRL)*, 2023. 3

A. Environments

A.1. Real-world Deployment

Our real robot employs a Unitree G1 platform, with an onboard Jetson Orin NX acting as the primary computing and communication device. The control policy receives motion-tracking target information as input, computes the desired joint positions for each motor, and sends commands to the robot’s low-level interface. The policy’s inference frequency is set at 50 Hz. The commands are sent with a delay kept between 18 and 30 milliseconds. The low-level interface operates at a frequency of 500 Hz, ensuring smooth real-time control. The communication between the control policy and the low-level interface is realized through LCM (Lightweight Communications and Marshalling)[27].

A.2. Simulation Tasks

In this section, we provide detailed state space information for the policy training.

Robot Proprioceptive States The robot proprioceptive states for the teacher and the student policy can be found in Table 6. Note that the student policy is trained on longer history length compared to the teacher, as it cannot observe privileged information but have to learn from a longer sequence of past observations.

Privileged Information The teacher policy leverages privileged information to obtain accurate motion-tracking performance. The complete information about the privileged states is listed in Table 7.

Tracking Target Information Both the teacher policy and student policy also take the motion tracking goal as part of their observations, which consists of the keypoint positions, DoF (joint) positions, as well as root movement information. The detailed components of the motion tracking target can be found in Table 8.

Reward Design In the main paper, we partially introduced our tracking-based reward design. Additionally, our reward also contains other penalties and regularization terms. The regularization reward components and their weights to compute the final rewards are introduced in Table 5. The final reward combines both the regularization with the tracking-based reward to train a robust RL policy.

Action Space The action is the target position of joint proportional derivative (PD) controllers, which is 23 dimensions for Unitree G1 and 21 dimensions for Unitree H1. Note that we are using an upgraded version of the Unitree H1. The older H1 has 19 DOF, while the new H1 has 21 DOF.

B. Model and Training Details

B.1. Baseline Implementation

Exbody[9]: The implementation of Exbody on the Unitree H1 robot is fully consistent with the original Exbody design, tracking only the upper-body keypoints and joint positions. For the Unitree G1 robot, we extended Exbody to tracking the all three dofs of the waist, while keeping other aspects identical. The key differences between Exbody and our method are as follows: Exbody focuses solely on upper-body motion tracking, does not utilize a teacher-student structure, uses a history length of only 5, and performs tracking entirely with local keypoints.

Exbody[†]: Exbody[†] is the full-body version of Exbody. It maintains most aspects of the original Exbody design but tracks the entire body’s keypoints and joint positions instead of just the upper body.

OmniH2O*[23]: The main difference between OmniH2O* and our method lies in the training phase. Specifically, OmniH2O* does not use the robot’s velocity as privileged information and relies solely on global tracking during training. For fairness, while OmniH2O* retains its original training method, we adapted it during testing to use local keypoints for evaluating tracking accuracy. Apart from this, we ensured that the observation space and reward design were consistent with the original OmniH2O implementation.

B.2. CVAE Architecture

Our CVAE model contains a transformer-based encoder and decoder. The encoder leverages bidirectional attention to extract the latent posterior variable z based on past M -step motions $m_{t-M:t}$ and future H -step motions $m_{t+1:t+H+1}$, and the decoder uses causal attention to predict future H -step motions $m_{t+1:t+H+1}$ based on z and $m_{t-M:t}$. Here m_t includes the current keypoint positions, root velocity, root angular velocity, position, and rotation.

During training, the transformer encoder f_{enc} takes $m_{t-M:t+H+1}$ and a classification token [CLS] and gets the output sequence. We adopt the corresponding output from the [CLS] token to obtain a posterior Gaussian distribution. A latent variable z is sampled and fed into the decoder alongside $m_{t-M:t+H+1}$ to generate $m_{t+1:t+H+1}$. During inference, z is set to the mean of the prior (i.e. zero). The CVAE is trained on the AMASS CMU motion capture dataset, conditioned on the past 50 frames of motions to predict the future 15 motions. We leverage a similar architecture and temporal ensembling strategy as proposed in [79] to obtain the predicted future motion. We provide the training and architecture hyper-parameters of CVAE in Table 9. During deployment, the humanoid robot takes the predicted motion from the CVAE and tracks it seamlessly in the real world.

Term	Expression	Weight
DoF position limits	$\mathbb{1}(d_t \notin [q_{\min}, q_{\max}])$	-10
DoF acceleration	$\ \ddot{d}_t\ _2^2$	$-3e^{-7}$
DoF error	$\ d_t - d_0\ _2^2$	-0.5
Energy	$\ \tau_t \dot{d}_t\ _2^2$	-0.001
Linear velocity (z)	$\ v_t^{\text{lin-z}}\ _2^2$	-1
Angular velocity (xy)	$\ v_t^{\text{ang-xy}}\ _2^2$	-0.4
Action rate	$\ a_t - a_{t-1}\ _2^2$	-0.1
Feet air time	$T_{\text{air}} - 0.5$	10
Feet velocity	$\ v_{\text{feet}}\ _1$	-0.1
Feet contact force	$\ F_{\text{feet}}\ _2^2$	-0.003
Stumble	$\mathbb{1}(F_{\text{feet}}^x > 5 \times F_{\text{feet}}^z)$	-2
Hip pos error	$\ p_t^{\text{hip}} - p_0^{\text{hip}}\ _2^2$	-0.2
Waist roll pitch error	$\ p_t^{\text{wrp}} - p_0^{\text{wrp}}\ _2^2$	-1
Ankle Action	$\ a_t^{\text{ankle}}\ _2^2$	-0.1

Table 5. Regularization rewards for preventing undesired behaviors for sim-to-real transfer and refined motion.

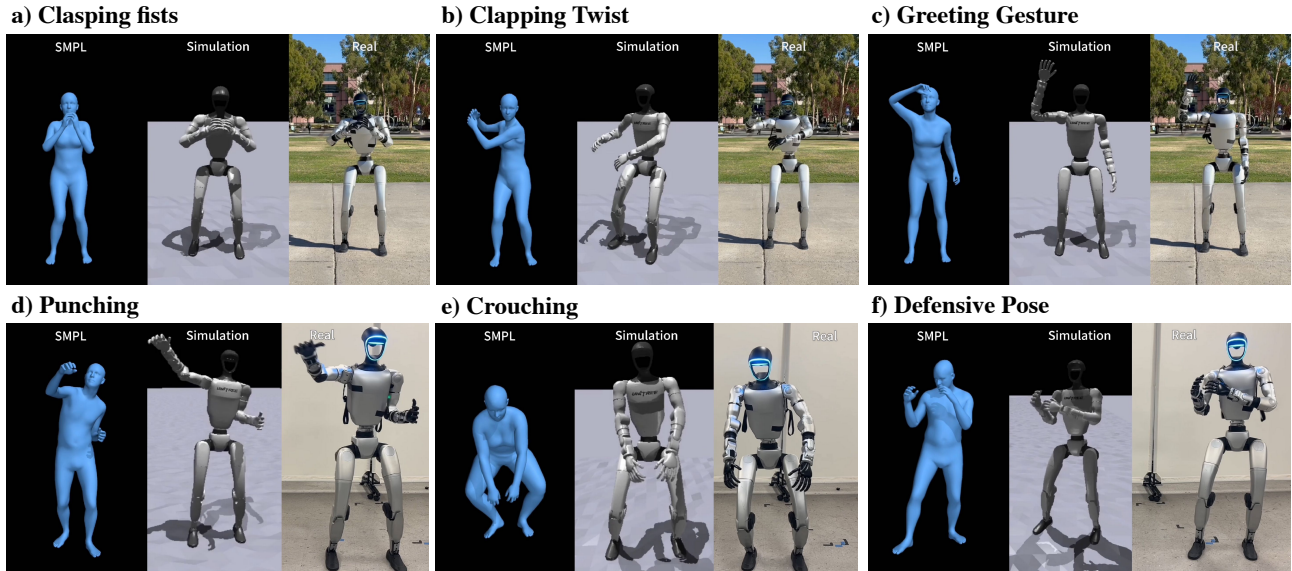


Figure 4. Sim-to-real experiment results showcasing diverse motions across SMPL, simulation, and real-world environments. Examples include: (a) Claspng Fists, (b) Clapping Twist, (c) Greeting Gesture, (d) Punching, (e) Crouching, and (f) Defensive Pose.

B.3. Policy Training Hyper-parameters

ExBody2 adopts a teacher-student training framework. The teacher policy is trained with standard PPO [60] algorithm on privileged information, tracking target and proprioceptive states. The student policy is trained with Dagger [59] without privileged information, but using longer history. For both teacher and student policies, we concatenate the corresponding inputs and feed them into MLP layers for policy learning. We provide the detailed training hyper-parameters for our teacher and student policy in Table 10.

C. More results of ExBody2

Figure 4 and Figure 5 provide illustrative examples of ExBody2’s motion tracking quality. Figure 4 highlights representative correspondences during specific motion tracking processes, showcasing the alignment between the original human SMPL model, simulation results, and real-world tracking outputs, which demonstrate high consistency. Additionally, Figure 5 presents a complete motion sequence, illustrating the correspondences as the motion progresses over time.




Figure 5. Sim-to-real experiment results showcasing striding walk motions across SMPL, simulation, and real-world environments.

State	Dimensions
DoF position (G1)	23
DoF position (H1)	21
DoF velocity (G1)	23
DoF velocity (H1)	21
Last Action (G1)	23
Last Action (H1)	21
Root Angular Velocity	3
Roll	1
Pitch	1
Yaw	1
Total Dim (G1)	75*25
Total Dim (H1)	69*25

Table 6. Proprioceptive states used in ExBody2. The rotation information is from IMU. 25 is the length of the history proprioception

State	Dimensions
DoF Difference (G1)	23
DoF Difference (H1)	21
Keybody Difference	36
Root velocity	3
Root angular velocity	3
Total dim (G1)	65
Total dim (H1)	63

Table 7. Privileged information used in ExBody2.

State	Dimensions
DoF position (G1)	23
DoF position (H1)	21
Keypoint position	36
Root Velocity	3
Root Angular Velocity	3
Roll	1
Pitch	1
Yaw	1
Height	1
Total dim (G1)	69
Total dim (H1)	67

Table 8. Reference information used in ExBody2.

Hyperparameter	Value
KL Loss Weight α	0.5
Smoothness Loss Weight β	0.5
Encoder Attention Layers	4
Decoder Attention Layers	4
Attention Heads	8
Transformer Dimension	512
Latent z Dimension	128
Context Length	50
Motion Prediction Horizon	15
Optimizer	AdamW
β_1, β_2	0.9, 0.999
Learning Rate	3e-4
Warmup Steps	2000
Weight Decay	0.01
Learning Rate Scheduler	cosine
Batch Size	512
Total Gradient Steps	60000

Table 9. Hyperparameters related to CVAE’s architecture and training.

Hyperparameter	Value
Optimizer	Adam
β_1, β_2	0.9, 0.999
Learning Rate	$1e^{-4}$
Batch Size	4096
Teacher Policy	
Discount factor (γ)	0.99
Clip Param	0.2
Entropy Coef	0.005
Max Gradient Norm	1
Learning Epoches	5
Mini Batches	4
Value Loss Coef	1
Entropy Coef	0.005
Value MLP Size	[512, 256, 128]
Actor MLP Size	[512, 256, 128]
Student Policy	
Student Policy MLP Size	[1024, 1024, 512]

Table 10. Hyperparameters related to the teacher and student policy’s training.